THE ADD-BACK · EPISODE 01

The Savings In The Model

Arc 1 · Pre-IC / confirmatory diligence

In the room: Deal team, investment committee

The question: Do I underwrite the automation savings in this model?

THE POSITION

SIGHTING

A target automated a back-office process and put $400,000 of run-rate savings into adjusted EBITDA. Three heads out, documented, termination letters on file.

THE CONSENSUS READ

Supportable. The headcount schedule ties, the reduction is real, annualise it.

THE MECHANISM

The reduction is real and the saving may not be. Automating an exception-rich process relocates work rather than removing it — out of a role with an owner and a cost centre, into a queue with neither. The heads left the schedule; the work went to a channel nobody codes to a department, so it never appears as an offsetting cost.

THE EXPOSURE

At 9x, that add-back is $3.6M of enterprise value resting on a cost search nobody ran.

THE TEST

Ask for eight weeks of message volume in whatever channel the failures land in. Rising, the saving is overstated. Flat, underwrite it and move on.

THE BRIDGE — CLAIMED TO REALISED

The case below is a composite built to the shape of a real quote-to-cash automation. Cited research figures are real and sourced; the two are kept separate throughout. A composite presented as a deal is the fastest way to lose this room.

LINE	MOVES	RUNNING
Claimed — $400,000 run-rate, three FTEs, back-office automation, month 11		**$400,000**
1. The headcount schedule is true and insufficient [changes the story]	-$0	$400,000
A QoE provider tests an add-back on four criteria: realised or contractually committed, documented, correctly characterised as recurring, and not offset by incremental…		
2. The exception queue	-$180,000 to -$260,000	$140,000 to $220,000
Exceptions ran ~40/week at go-live and ~190 by month 11, in a channel nobody had plotted.		
3. The distribution split, and what it does to the KPI [changes the story]	-$0	$140,000 to $220,000
The composite reported a 60% reduction in median cycle time.		
4. Where I was wrong — I put the boundary where the work looked different [where I was wrong]	-$0	$140,000 to $220,000
Given a queue with no owner, my recommendation was the obvious one: classify exception-prone transactions at intake and route them to a named person with an SLA. Cost: Operators stopped trusting the routing and quietly began checking everything, erasing the gain on the clean majority too.		
5. What survives [residual]	+$0	$140,000 to $220,000
Note what the bridge did not conclude.		
Realised		**$140,000 to $220,000**

THE DOWNSIDE

Multiple: 9x, illustrative

A $400,000 adjustment at an illustrative 9x is $3.6M of purchase price. If the offsetting-cost search lands where the composite lands, $180,000–$260,000 of it is unsupported. Substitute your own comps.

$1.6M to $2.3M of enterprise value — roughly 0.4 to 0.6 turns on the adjustment itself

When it surfaces. Pre-IC it is a price negotiation and costs you nothing but time. Discovered in the buyer's confirmatory diligence at exit it is a retrade after LOI, at the point of maximum asymmetry, against a seller who has already told their LPs the process is closing. The same finding is worth a discount going in and a discount coming out — and one of those is yours.

The Savings In The Model

The question: Do I underwrite the automation savings in this model?

Paste into a request list or a management agenda. Each question resolves to an artifact, not to a characterisation.

THE DILIGENCE PACK

1. Export the full message history, with timestamps, of every channel, shared inbox or ticket queue that receives failed or exception transactions from the automated process, for the twelve months pre- and post-implementation.

Artifact: Raw export. Not a summary.

2. Provide a headcount schedule for the three eliminated roles, and separately, the current time allocation of every named individual appearing more than twenty times in that export.

Artifact: Schedule plus named list.

This is the offsetting-cost search, and it is the question that moves the number.

3. Provide cycle time for the process at median, mean and 90th percentile, monthly, for twenty-four months.

Artifact: Transaction-level data, not a dashboard screenshot.

4. Identify every transaction type requiring a data element not present in the transaction's own record, and for each, describe how that element was obtained before implementation.

Artifact: Written response from operations, not IT.

The answer is frequently a person's name.

5. Provide any remediation proposal, build estimate or vendor scope prepared in response to exception volume, and the board or management materials that reviewed it.

Artifact: The documents themselves.

If a remediation is already scoped, the exception problem is known internally and unpriced in the model.

DISQUALIFIER

If the response to question 1 is that no such channel exists, that is not a clean answer. Ask where failures go instead. There is always somewhere.

SOURCES — EVERY MEASUREMENT WITH ITS SAMPLE

CLAIM	SOURCE	SAMPLE / CLASS
Quote-to-cash automation case: $400,000 claimed, exceptions 40/wk rising to 190/wk by month 11, 1.5–2.0 FTE absorbed across adjacent functions	Composite, built to the shape of a real engagement	n/a — composite (composite)
QoE four-criteria test for an add-back: realised or committed, documented, correctly characterised, not offset by incremental costs	Standard quality-of-earnings practice	n/a — practice standard (reported)
9x entry multiple used throughout	Illustrative mid-market comp	n/a — illustrative, arithmetic exposed (illustrative)

One standing caveat. Every number in this show is somebody else's measurement, and I'll tell you whose, with the sample. None of it is diligence on your deal. Do that yourself.